

FINNEGAN, HENDERSON, FARABOW, GARRETT & DUNNER, L.L.P.

1300 I STREET, N. W.
WASHINGTON, DC 20005-3315

202 • 408 • 4000
FACSIMILE 202 • 408 • 4400



ATLANTA
404 • 653 • 6400
PALO ALTO
650 • 849 • 6600

WRITER'S DIRECT DIAL NUMBER:

TOKYO
011 • 813 • 3431 • 6943
BRUSSELS
011 • 322 • 646 • 0353

(202) 408-4039

June 12, 2000

ATTORNEY DOCKET NO. 07030.0017-00

Box PATENT APPLICATION
Assistant Commissioner for Patents
Washington, D.C. 20231

Re: New U.S. Patent Application
Title: PROCESSOR FOR NETWORK COMMUNICATIONS
WITH EFFICIENT USE OF PROCESSOR CYCLES
Inventor: Robert Gelinis et al.

Sir:

We enclose the following papers for filing in the United States Patent and Trademark Office in connection with the above patent application.

1. Application - 20 pages, including 6 independent claims and 26 claims total.
2. Drawings - 6 sheets of informal drawings (Figures 1 - 6).
3. A check for \$1,032.00 representing a \$690.00 filing fee and \$342.00 for additional claims.

This application is being filed under the provisions of 37 C.F.R. § 1.53(f). Applicants await notification from the Patent and Trademark Office of the time set for filing the Declaration.

Please accord this application a serial number and filing date.



09591510 06/12/00

FINNEGAN, HENDERSON, FARABOW, GARRETT & DUNNER, L.L.P.

The Commissioner is hereby authorized to charge any additional filing fees due and any other fees due under 37 C.F.R. § 1.16 or § 1.17 during the pendency of this application to our Deposit Account No. 06-0916.

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

By: E. Robert Yoches
E. Robert Yoches
Reg. No. 30,120

ERY:bbp
Enclosures

UNITED STATES PATENT APPLICATION

OF

ROBERT GELINAS

W. PATRICK HAYS

SOL KATZMAN

AND

WILLIAM J. DALLY

FOR

PROCESSOR FOR NETWORK COMMUNICATIONS WITH EFFICIENT USE OF

PROCESSOR CYCLES

0022190" 01916960

I. BACKGROUND OF THE INVENTION

This invention relates generally to the field of network communication processors, and more specifically to the field of system device instructions and context switching.

Network communication systems demand real-time performance. The performance of conventional processors in network communication systems is degraded by long latency accesses, especially to shared resources. For example, in order to look up data in a table lookup unit, a processor must send an operation with data to the table lookup unit (TLU) commanding the TLU to look up data in a table. After performing the lookup operation, the TLU stores the resulting data internally. The processor sends a load command requesting that the TLU load the result on the bus and return the data to the processor. This procedure requires two bus transactions initiated by the processor. Therefore, it would be desirable to have a single transaction both command the device to perform an operation and provide the result to the processor.

Another latency problem is that some conventional processors will await receipt of the result of the look up before processing other instructions. One way of dealing with this problem is to perform instructions in a different thread while a first thread awaits data. This is called a context switch. Context switches performed in software, store all data in the processor registers in memory and then use the processor registers for a new context. This requirement to store and restore data using a single set of registers wastes processor cycles. Therefore, it would be desirable to have a context switch performed that does not waste processor cycles.

II. SUMMARY OF THE INVENTION

Systems and methods consistent with the present invention allow for performing a single transaction that supplies data to a device and commands the device to perform an action and return the result to a processor.

In addition, systems and methods consistent with the present invention further allow for performing a context switch with no stall cycles by using an independent set of registers for each context.

A processing system consistent with the present invention includes a processor configured to formulate an instruction and data for sending to a device. The formulated instruction requests that the device perform a command and return data to the processor. A bus controller is configured to generate a system bus operation to send the formulated instruction and data along with a thread identifier to the device.

A processor consistent with the present invention executes instructions in threads. The processor includes a context register file having a separate set of general registers for a plurality of contexts, where the threads are each assigned a separate context, and context control registers having a separate set of control registers for the plurality of contexts.

Another processing system consistent with the present invention includes a processor configured to formulate an instruction and data, from a thread associated with a first context, for sending to a device, the instruction requesting the device to perform a command and return data to the processor, and perform a context switch to switch from processing the first context to a second context. A bus controller is configured to generate a system bus operation to send the formulated instruction and data along with a thread identifier to the device.

A method consistent with the present invention processes a single instruction that both requests a system device operation and requests the system device return data, the method comprising the steps of fetching an instruction from memory, forming a descriptor, constructing a system bus address, initiating a system bus operation to request a device to perform an operation and return data to a processor identified in a thread identifier, and retrieving return data from a system bus based on the thread identifier provided with the returned data.

Another method consistent with the present invention switches between contexts using a processor having a context register file having a separate set of general registers for a plurality of contexts, each set of registers being associated with a thread, and context control registers having a separate set of control registers for the plurality of contexts, the method comprising the steps of receiving a context switch instruction, receiving an identifier of a next context to activate from the scheduler, performing a next instruction in a current context, and pointing a processor program counter to the context program counter in the context control register associated with the next context.

III. BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate one embodiment of the invention and, together with the description, serve to explain the objects, advantages, and principles of the invention. In the drawings:

Fig. 1 is a block diagram of a processing system consistent with methods and systems of the present invention;

Fig. 2a shows a context register file consistent with methods and systems of the present invention;

Fig. 2b shows a context control file consistent with methods and systems of the present invention;

Fig. 3 shows an instruction format consistent with methods and systems of the present invention;

Fig. 4 is a flowchart showing the steps for processing a write descriptor load word instruction consistent with methods and systems of the present invention;

Fig. 5 is a flowchart showing the steps of a method for processing a write descriptor load word with a context switch consistent with methods and systems of the present invention; and

Fig. 6 is a flowchart showing the steps of a method for completing the load word for the instruction in Fig. 5.

IV. DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Reference will now be made in detail to embodiments consistent with this invention that are illustrated in the accompanying drawings. The same reference numbers in different drawings generally refer to the same or like parts.

Processing systems for network communications require several bus and processor cycles to perform a write to a system device and a read from the system device. Systems and methods consistent with the present invention provide a single instruction that directs a device to read and load data when the device is ready. In accordance a preferred embodiment, the single instruction

includes a thread identifier so that the device can send the data back to the requesting thread at the processor.

In addition, systems and methods consistent with the present invention provide for a context switch that prevents the introduction of stall cycles by using a different set of registers for a plurality of threads. In this manner, processing can switch quickly from one set of registers used by one thread to a different set of registers used by another thread. As used herein, the term thread describes a set of program instructions or a software program that relies on a context register set to perform a particular task.

Fig. 1 shows an exemplary processing system that may be used in systems and methods consistent with the present invention. Processor 100 is preferably a RISC type processor that may include, among other elements, those in Lexra's LX4180 processor. In this example, processor 100 connects to instruction memory 120, which may be, for example, a cache, RAM or DRAM.

Processor 100 includes a context register file 200 and context control registers 210. As used herein, a context is an independent set of general registers in context register file 200 and control registers in context control register 210 that are used in executing a thread. As stated, a thread may be software that relies on the contents of the context registers to perform a particular task. The term context may also generally be used to refer to a thread currently using the context's registers. Processor 100 further includes a processor program counter (PPC) 110 that points to the program counter of an active context stored in a context program counter within the context control registers 210.

parts, many of which are common and whose description is unnecessary to understand the systems and methods consistent with the present invention.

Fig. 2a shows an exemplary context register file 200 having 8 contexts, context 7 through context 0. In this figure, each context has 32 physical general registers, but the number of contexts and the number of registers may vary depending on the complexity of the particular system, the amount of data communication on the system bus, the number of system devices present, etc.

Fig. 2b shows an exemplary context control file 210 having 3 control registers for each of the 8 contexts shown in Fig. 2a. Context control file 210 includes a context program counter (CXPC) 212 for keeping track of the next instruction to be executed in the context and a context status register (CXSTATUS) 214 having a wait load bit, which, when set, indicates that the context is awaiting data from an external device. CXSTATUS 214 may include additional status information such as an indication that the context requires external events or data to complete its task. A write address register 216, also within context control file 210, is configured to store the address of a general purpose register in an inactive context that may be awaiting data from an external device.

Fig. 3 is an exemplary representation of an instruction 300 stored in instruction memory 120. Instruction 300 includes an opcode field 310 and sub-opcode field 360 that indicate the particular operation requested. The requested operations may be commands such as read, write, and write-split read. In this example, a write-split read is an instruction that writes to a system device and directs the device to return read data when available. Instruction 300 also includes rS 320, rT 330, and rD 340; fields referring to the general purpose registers in Fig. 2a. The

identified registers hold data used by the instruction or the registers that will ultimately be receiving the instruction results. In a write-split read instruction, for example, rS 320 and rT 330 identify the registers holding data that will be written to system device 160 at system device address 350. rD 340 is the identifier of the destination register, indicating the location in which the result of the load instruction should be stored. The identity of register rD may be stored in the write address register 216 so that when load data is returned, processor 100 reads the context control file 210 to determine the particular register in which to write the result.

Fig. 4 shows the steps of a method 400 for processing a write-split read instruction consistent with the methods and systems of the present invention. First, processor 100 fetches instruction 300 from instruction memory 120 based on a value in PPC 110 (step 410). Processor 100 then forms a 64 bit descriptor by concatenating bits [63:32] of register S 320 and bits [31:0] of register T 330 (step 420). Processor 100 constructs a system bus address using device address 350 provided in the instruction (step 430). The actual device address is less than 32 bits, so the remaining system bus address bits are set to zero or some constant predefined value.

Following the construction of the system bus address, processor 100 initiates a system bus operation to write the descriptor to the device, having the device perform some function, and requests that the device provide a read word response back to the processor identified with a GTID (step 440). Bus controller 140 sends out instruction 300 to the device address including data, the command, and a GTID. System device 160 saves the descriptor in a memory, performs an operation using information in the descriptor, and returns the result of the operation as read data directed to the processor identified in the GTID (step 450). Bus controller 140 then receives

a read word or twin word response from the system device (step 460). Finally, processor 100 writes the received data to rD register 340 (step 470) thus, completing the operation.

Fig. 5 show the steps of a method 500 for processing a write descriptor load word (WDLW) instruction in accordance with systems and methods of the present invention.

Referring to Fig. 5, Processor 100 initially fetches instruction 300 from instruction memory 120 based on the value in PPC 110 (step 510). Using this value, processor 100 forms a 64-bit descriptor by concatenating bits [63:32] of register S 320 and bits [31:0] of register T 330 (step 520). Processor 100 next sets the wait load bit in context status register 210 of the active context (step 530). Processor 100 then constructs a system bus address using device address 350 provided in the instruction (step 540). The device address is less than 32 bits, so the remaining system bus address bits are set to zero or some constant predefined value.

Once the system bus address is constructed, processor 100 initiates a system bus operation to write the descriptor to the device and requests that the device provide a read word response (step 550). Processor 100 stores the register identified in rD 340 in write address register 216 in the active context's control file 210 indicating the register that will receive any returned data from system device 160 (step 560).

Steps 565-590 describe the steps used to perform a context switch in systems and methods consistent with the present invention. Processor 100 first receives an identifier of the next context to be activated from scheduler 130 (step 565). Processor 100 then performs the following instruction in the active context (step 570). By performing the next step in this instruction before moving on to the next context, the processor is able to execute an instruction, and is performing useful work instead of stalling for a cycle while the context switch is

performed. Processor 100 then stores program counter (PC) of the next instruction in this active context in the CXPC 212 of the active context (step 580). Processor next points PPC to CXPC 212 of the new context designated by scheduler 130 (step 590).

Fig. 6 shows the remaining steps 600 for completing the load word portion of the WDLW instruction described in the method of Fig. 5. After system device 160 receives the command, data, and the GTID from system bus 150, it writes the descriptor to a memory. System device 160 then performs any requested function and loads the resulting data onto system bus 150 along with the GTID (step 610). Upon receiving the read word response from system bus 150 (step 620), bus controller 140 forwards it to processor 100. Processor 100 writes this read word to the register indicated in the write address register 216 by obtaining the identity of the originating context from the GTID (step 630). Processor 100 next clears the originating context's wait load flag in CXSTATUS register 214, indicating that the context is available for execution (step 640). Finally, scheduler 130 monitors the wait load flags of all of the contexts and will select this context when appropriate (step 650).

There are many variations that may be made consistent with the present invention. For example, in another embodiment, system device 160 returns a twin word in response to a write twin word read twin word instruction (WDLT). Further, while the implementations above specifically mention word or twin word data reads and writes, systems and methods consistent with the present invention may be used with other sized data reads and writes. In addition, there may be multiple processors sharing the system bus and accessing the system bus devices.

The foregoing description is presented for purposes of illustration and description. It is not exhaustive and does not limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practicing the invention. The scope of the invention is defined by the claims and their equivalents.

09591510-061200

WHAT IS CLAIMED IS:

1. A processing system comprising:
a processor configured to
formulate an instruction and data for sending to a device, said instruction requesting the device to perform a command and return data to the processor; and
5 a bus controller configured to
generate a system bus operation to send the formulated instruction and data along with a thread identifier to the device.
2. The processing system of claim 1, wherein the device is a table lookup unit.
3. The processing system of claim 1, wherein the bus controller is further configured to receive the device return data from a system bus along with the thread identifier.
4. A processor that executes instructions in threads, said processor comprising:
a context register file having a separate set of general registers for a plurality of contexts, wherein said threads are each assigned a separate context; and
context control registers having a separate set of control registers for the plurality of contexts.
5. The processor of claim 4, wherein the context register file includes 32 general registers for 8 contexts.

[illegible]

a context status register; and

a write address register.

7. The processor of claim 6, wherein the context program counter holds a program counter pointing to a next instruction in an associated context.

8. The processor of claim 6, wherein the context status register holds data that indicates whether an associated context is awaiting data from an external source.

9. The processor of claim 6, wherein the write address register stores an address of a register associated with a context that is awaiting data.

10. The processor of claim 6 further including:
a scheduler configured to select a context to activate.

11. The processor of claim 10 further including:

means for receiving a context switch instruction;

means for receiving an identifier of a next context to activate from the scheduler;

means for performing a next instruction in a current context; and

5 means for pointing a processor program counter to the context program counter in the context control register associated with the next context.

12. A processing system comprising:

a processor configured to

formulate an instruction and data, from a thread associated with a first context, for sending to a device, said instruction requesting the device to perform a command and return data to the processor;

perform a context switch to switch from processing the first context to a second context; and

a bus controller configured to

10 generate a system bus operation to send the formulated instruction and data along with a thread identifier to the device.

13. The processing system of claim 12, wherein the processor is further configured to store an address of a destination register associated with the first context receiving the return data.

14. The processing system of claim 13, wherein the bus controller is further configured to receive the device return data from a system bus along with the thread identifier.

15. The processing system of claim 14, wherein said processor is configured to write the return data to the stored destination register associated with the first context; and
clear a wait bit in a status register associated with the first context.

16. The processing system of claim 12, wherein the device is a table lookup unit.

17. The processing system of claim 12 said processor further including:
a context register file having a separate set of general registers for the first and second contexts; and
context control registers having a separate set of control registers for the first and second contexts.

18. The processing system of claim 17, wherein the context register file includes 32 general registers for the each context.

19. The processing system of claim 17, wherein the context register file includes for each context:

- a context program counter;
- a context status register; and
- a write address register.

20. The processing system of claim 19, wherein the context program counter holds a program counter pointing to a next instruction in an associated context.

21. The processing system of claim 19, wherein the context status register holds data that indicates whether an associated context is awaiting data from an external source.

22. The processing system of claim 19, wherein the write address register stores an address of a register associated with a context awaiting data.

23. The processing system of claim 19 further including:
a scheduler configured to select the second context to activate.

24. The processing system of claim 23 further including
means for receiving an identifier of the second context to activate from the scheduler;
means for performing a next instruction in the first context; and
means for pointing a processor program counter to the context program counter in the
5 context control register associated with the second context.

25. A method for processing a single instruction that both requests a system device
operation and requests the system device return data, said method comprising the steps of:
fetching an instruction from memory;
forming a descriptor;
constructing a system bus address;
initiating a system bus operation to request a device to perform an operation and return
data to a processor identified in a thread identifier; and
retrieving return data from a system bus based on the thread identifier provided with the
returned data.

26. A method for switching between contexts using a processor having a context
register file having a separate set of general registers for a plurality of contexts, each et of
registers associated with a thread, and context control registers having a separate set of control
registers for the plurality of contexts, said method comprising the steps of:
5 receiving a context switch instruction;
receiving an identifier of a next context to activate from the scheduler;

performing a next instruction in a current context; and

pointing a processor program counter to the context program counter in the context control register associated with the next context.

00000000 00000000 00000000 00000000

ABSTRACT

Systems and methods that allow for performing a single transaction that both commands a device to perform an action and return the result to a processor without the processor having to send a separate request for the result. In addition, a processor may perform a context switch switching between threads using separate context registers. A context switch allows a processor to switch from processing a thread that is waiting for data to one that is ready for additional processing.

09594510-051200

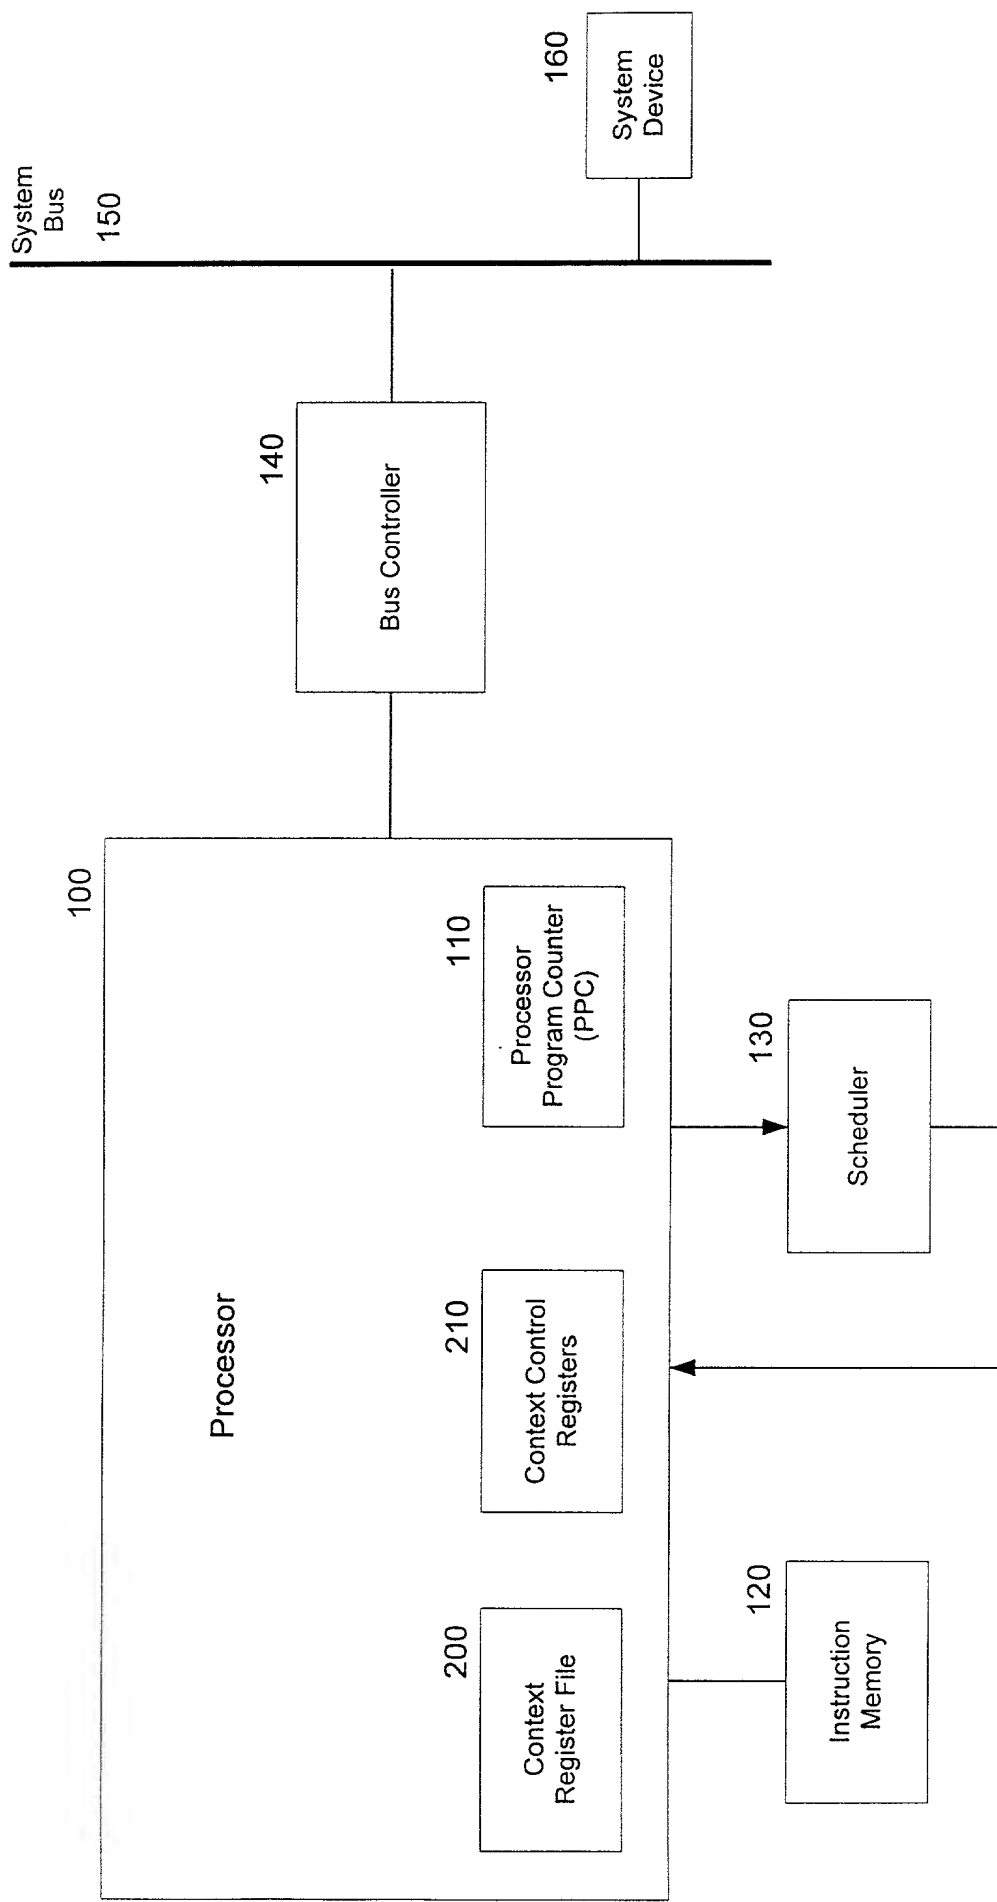


Fig. 1

200




Fig. 2a

210




Fig. 2b

00591510-061200




Fig. 3

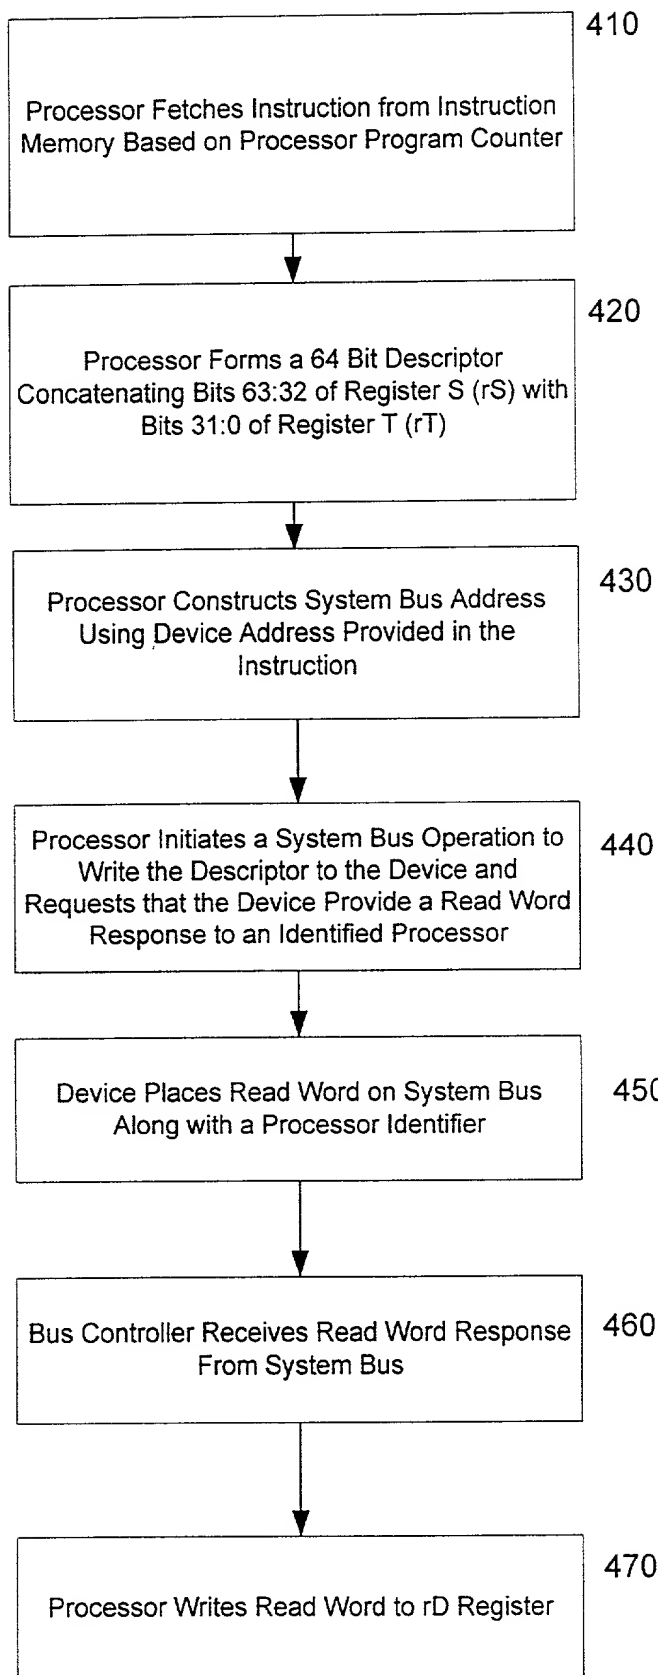


Fig. 4



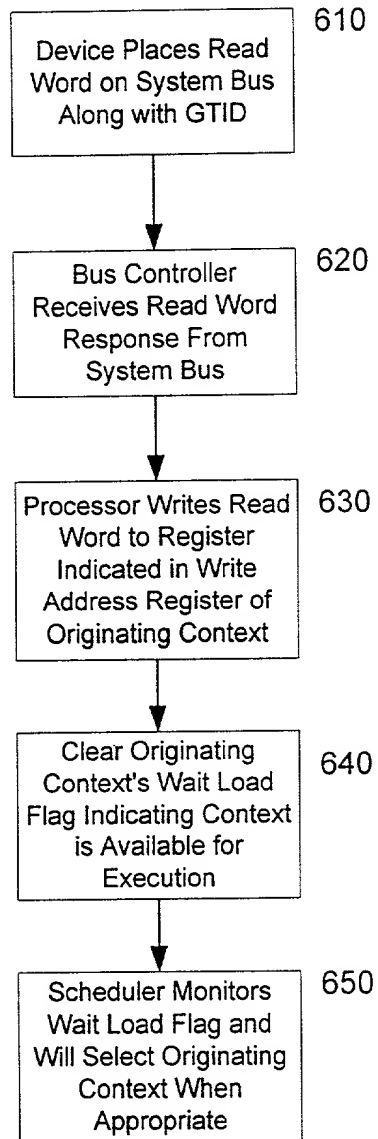


Fig. 6